Interval Arithmetic and Interval-Aware Operators for Genetic Programming*

Grant Dick
Department of Information Science
University of Otago
Dunedin, New Zealand
grant.dick@otago.ac.nz

arXiv:1704.04998v1 [cs.NE] 17 Apr 2017

ABSTRACT

Symbolic regression via genetic programming is a flexible approach to machine learning that does not require up-front specification of model structure. However, traditional approaches to symbolic regression require the use of protected operators, which can lead to perverse model characteristics and poor generalisation. In this paper, we revisit interval arithmetic as one possible solution to allow genetic programming to perform regression using unprotected operators. Using standard benchmarks, we show that using interval arithmetic within model evaluation does not prevent invalid solutions from entering the population, meaning that search performance remains compromised. We extend the basic interval arithmetic concept with 'safe' search operators that integrate interval information into their process, thereby greatly reducing the number of invalid solutions produced during search. The resulting algorithms are able to more effectively identify good models that generalise well to unseen data. We conclude with an analysis of the sensitivity of interval arithmetic-based operators with respect to the accuracy of the supplied input feature intervals.

CCS CONCEPTS

•**Computing methodologies** → **Genetic programming; Model verification and validation;** Supervised learning;

KEYWORDS

genetic programming; symbolic regression; interval arithmetic

1 INTRODUCTION

For over two decades, researchers have explored the use of genetic programming (GP) to evolve models in a regression setting [12]. The resulting symbolic regression approach is interesting as it simultaneously searches for a suitable model structure and its corresponding parameters. This presents a highly desirable framework that frees the user from a priori decisions pertaining to model structure, allowing novel and potentially insightful models to be discovered and provide better understanding of the problem.

A long-known issue with symbolic regression is the need for protected operators [12]. Because we do not know up-front the conditions under which a particular operation will be applied, we need to ensure that its application will produce results that will not corrupt any subsequent operations in the model. By far, the most common solution to this is to develop customised 'protected' operators with built-in exception handling, such as for division by zero. However, the very nature of GP is to exploit this behaviour to create novel solutions that fit well to training data at the expense of generalisation performance. An alternative solution that has been proposed is to identify the ranges of input features and apply interval arithmetic to identify safe and valid use of operators. These intervals ideally capture knowledge of the problem domain that can help identify whether or not a particular operation will remain valid over all possible inputs. If interval arithmetic identifies an operation as producing an undefined interval on some inputs, then we can flag the individual containing that operation with a poor fitness to discourage its selection and propagation.

Previous work has demonstrated some success in applying interval arithmetic over protected operators [8]. However, most work incorporating interval arithmetic into GP uses it purely in an evaluative framework — the dynamics of interval arithmetic within the population, and how it can be further exploited within GP, remain largely unexplored. The goal of this paper is to explore how interval arithmetic may be used to change the behaviour of GP by using interval arithmetic to guide safe and effective use of search operators. In doing so, it is shown that interval arithmetic, when used purely as an evaluative tool, cannot prevent invalid individuals appearing in the population during run. It is also shown that relying on selection pressure to eliminate invalid solutions from the population is ineffective, as standard mutation and crossover operators frequently generate offspring that produce undefined execution intervals. New search operators are presented in this paper that attempt to honour the intervals presented by the problem domain to ensure that offspring remain valid during search. This increases the rate of search within the population, and leads to more rapid evolution towards fit models. In addition to these new operators, the paper also examines how effectively input intervals can be estimated from available data, and to what extent GP using interval arithmetic is sensitive to the quality of these estimated intervals.

The remainder of this paper is structured as follows: §2 provides an introduction to interval arithmetic and related work in GP; §3 demonstrates the dynamics of GP using interval arithmetic within evaluation, and informs the process of developing new operators presented in §4; the new operators are examined using several benchmark problems §5, and the sensitivity of interval arithmetic-based GP using estimated input intervals is also examined; finally §6 concludes the paper and suggests areas of future work.

2 OPERATOR SAFETY IN GP

Identification of models through symbolic regression is one of the most thoroughly explored areas of GP research — it is suggested that over a third of GP research either uses symbolic regression

*An abridged version of this paper appears in [6]

Table 1: Interval calculations for some of the operators used in this paper.

Operator	Resulting Interval
$[a,b]+[c,d]$	$[a+c, b+d]$
$[a,b]-[c,d]$	$[a-d, b-c]$
$[a,b]\times[c,d]$	$[min(a\times b, a\times d, c\times b, c\times d), max(a\times b, a\times d, c\times b, c\times d)]$
$[a,b]\div[c,d]$	$\begin{cases}\text{undefined} & \text{, if } 0 \in [c,d] \\ [min(a\div b, a\div d, c\div b, c\div d), \\ \ max(a\div b, a\div d, c\div b, c\div d)] & \text{, otherwise}\end{cases}$
$e^{[a,b]}$	$[e^a, e^b]$
$log_e[a,b]$	$\begin{cases}\text{undefined} & \text{, if } a \le 0 \\ [log_e a, log_e b] & \text{, otherwise}\end{cases}$

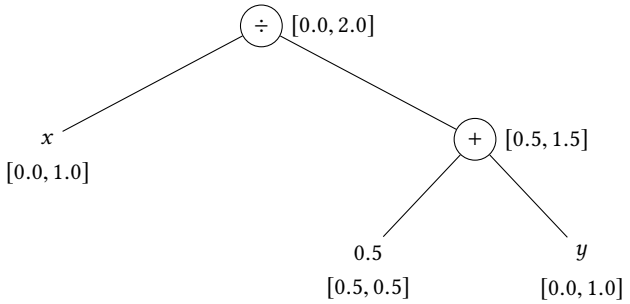

Figure 1: Static analysis of a simple parse tree. Interval information percolates up the known intervals of the terminal nodes to the root of the tree.

problems as benchmarks, or actively explores new ways in which to conduct symbolic regression [17]. Unlike traditional machine learning approaches to regression, symbolic regression through GP does not require the up-front selection of model structure. Instead, the GP system is free to explore a range of possible model structures, giving rise to the possibility of useful insights being discovered through evolution. However, to achieve this flexibility, the GP must provide closure, where the result of any operation must be valid as a possible input to subsequent operations. This presents a problem to some operators, such as division or computing the square root of a number. When operators are used that do not define valid outputs over all possible inputs, the GP system must incorporate some logic to handle these exceptional cases.

By far, the most common solution to the closure problem in symbolic regression is to adopt protected versions of standard mathematical operators. For example, the standard division operator can be replaced by a protected version where any division by zero is replaced with the constant value 1. This is the solution presented by Koza in his early work exploring GP [12]. While simple in nature and implementation, the particular nuances of protected operators are often exploited by GP to create solutions that fit well to training data, but lead to poor generalisation. As an alternative to using protected operators, Keijzer pioneered the idea of using information about the known ranges of input variables to compute the expected output interval of an evolved model [8]. If we know the interval of a variable (or equivalently a sub-expression), then we can use simple rules to compute the expected interval of a mathematical operation. Some examples of this are shown in Table 1. Following this, we can chain these operations from the known intervals of our input features through to the root of a given parse tree. An example of this is given in Figure 1: starting with the known intervals of our terminals x and y (in this case, both defined over $[0, 1]$), and our constant 0.5, we can work out the interval of the addition operator, and then finally the division at the root of the tree.

Since Keijzer's original paper on using interval arithmetic in GP, interval arithmetic has seen use primarily as a tool to evaluate individual solutions and assign fitness penalties to individuals that present potentially invalid execution intervals [9, 15, 16]. Kotanchek et al. use interval arithmetic to help select intervals with robust intervals as part of a multi-objective symbolic regression framework [11]. Others have extended the spirit of interval arithmetic into using other forms of validation, such as affine arithmetic [13]. However, the influence of interval arithmetic within the search process itself remains a largely unexplored concept.

3 BEHAVIOUR OF INTERVAL ARITHMETIC

This paper aims to uncover some of the within-run properties of interval arithmetic in GP with the aim of developing new operators that exploit interval information to provide more effective search. We start by exploring a simple problem from previous work [8]. This function, referred to as *Keijzer-10* is defined over two variables:

$$f(x_1, x_2) = x_1^{x_2} \tag{1}$$

where x_1 and x_2 are both defined over the interval $[0, 1]$. In previous work, this problem proved a challenge for GP, where without interval arithmetic, GP failed to find a meaningful solution in 98% of runs. Here we will explore this problem again, using a standard implementation of GP, once with protected operators and then again with unprotected operators, and then an implementation augmented with interval arithmetic to perform static evaluation. The static evaluation was performed as follows: after creation through initialisation, mutation or crossover, the known intervals of the terminal nodes were supplied to the tree, and were recursively passed through the tree to compute the solution's execution interval. If the interval is deemed valid, then the individual is subsequently passed on to a normal evaluation process. If the interval is deemed invalid, then the individual is assigned the lowest possible fitness, with the intended effect being to remove it from selection in the next generation. The parameters for the GP system are defined in Table 2 and were settled upon after examining recent work and some small experimentation trying different population size and generations [3, 5]. As interval arithmetic is attempting to characterise aspects of generalisation performance, a hold-out set testing approach was used similar to that in previous work. For each run, 20 instances were sampled uniformly from the problem domain and used as a training set. For testing, a mesh sample over the problem domain was used to generate 10000 points uniformly over the problem space. In total, 100 separate runs were performed, each using a different training set. For analysis, we explored three aspects, the training performance (in terms of the best individual in the population at a given generation), the performance of the

Table 2: Parameters used for all experiments in this paper.

Parameter	Setting
Population size	200
Generations	250
Initialisation	Ramped Half-and-Half
Min. initial depth	2
Max. initial depth	6
Mutation	Subtree (max. depth 4)
Crossover	Subtree swap
Max. offspring depth	17
Crossover prob.	0.3
Mutation prob.	0.7
Selection	Tournament (size: 3)
Elitism	1 (fittest from previous generation)
Function set	$+, -, \times, \div$, sin, cos, exp, log
Terminal set	Input features: $x_1, x_2, \ldots, x_p$, plus ephemeral random constants from a uniform distribution over [-5, 5]

best individual on the test set, and the number of invalid individuals residing in the population in a given time frame. Error performance was recorded using the root-relative squared error (RRSE) measure:

$$RRSE(y, \hat{y}) = \sqrt{\frac{\sum_{i=1}^{|y|} (y_i - \hat{y}_i)^2}{\sum_{i=1}^{|y|} (y_i - \bar{y})^2}}, \quad (2)$$

where y and $\hat{y}$ are the recorded response values of the data set and the model predictions, respectively. The RRSE is analogous to the normalised root mean square error used in previous work. In all graphs presented in this paper, the statistic being plotted is the median, with corresponding shaded areas representing a 95% confidence interval of the median.

The results of the initial analysis using *Keijzer-10* are shown in Figures 2–4. The results are rather interesting: in terms of training performance, both protected and unprotected operators appear to allow the GP system to evolve at a faster rate than when using interval arithmetic and static analysis. A possible reason for this is provided by the results in Figure 3, which graphs the number of invalid individuals present in the population in a given generation.[1]. As can be seen, the use of interval arithmetic means that more offspring are identified as being invalid. Intuitively, this makes sense, as interval arithmetic provides a second chance at flagging an individual as problematic that may have been missed through the standard evaluation process. However, the use of interval arithmetic has the effect of reducing the size of the pool of individuals that selection can effectively work on to produce the next generation, essentially reducing the population size. It is interesting to note that the proportion of invalid solutions in the population does not substantially decrease over time.

While the integration of interval arithmetic into GP appears to compromise training performance, it is clearly offset by stronger generalisation performance. Examining Figures 4 and 5 show the

[1]We define an 'invalid' individual here as either one that has produced errors while being evaluated on the training data, or one that was identified as potentially producing errors as a consequence of static analysis using interval arithmetic

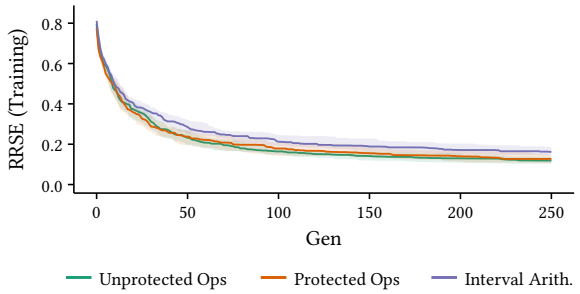

Figure 2: Training performance on the *Keijzer-10* problem.

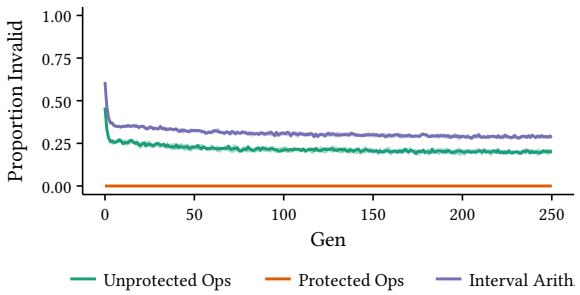

Figure 3: Proportion of population with invalid performance on the *Keijzer-10* problem.

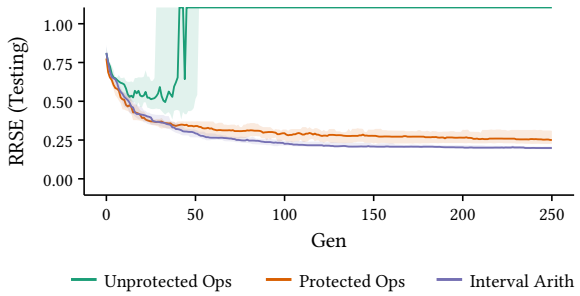

Figure 4: Test performance on the *Keijzer-10* problem.

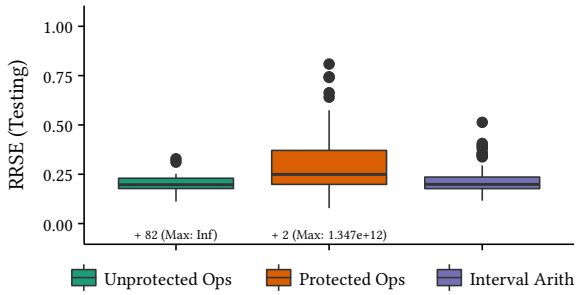

Figure 5: Test performance of solutions in the final generation on the *Keijzer-10* problem. The number under a given boxplot indicates the number of runs where an RRSE ¿ 1 was encountered (and the number in brackets indicates the largest RRSE encountered).

good generalisation performance provided by using interval arithmetic over protected and unprotected operators. In agreement with previous work, unprotected operators failed to find a meaningful solution to the problem in 82 out of 100 trials. Protected

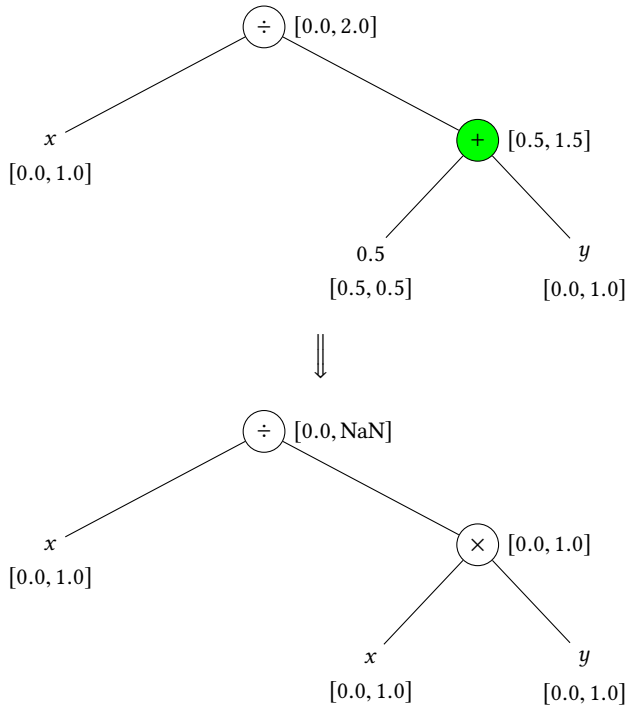

Figure 6: Example of mutation producing an offspring with an invalid output interval.

operators were also unable to properly manage this problem. It is interesting to note, however, that on several equations, GP using either protected or unprotected operators evolved a solution that either matched, or very nearly matched, the following equation:

$$f(x_1, x_2) = e^{\log_e(x_1)\cdot x_2}. \tag{3}$$

This particular equation is a perfect match for the *Keijzer-10* problem on all inputs except for when $x_1 = 0$, where it is undefined. Clearly, in these runs, there were no training instances sampled with x_1 set to zero. Finding such a solution was impossible for GP using interval arithmetic, regardless of training set conditions, as such a solution would be ruled out during static analysis.

4 INTERVAL-PRESERVING OPERATORS

The results presented in the previous section align with previous work in suggesting that including interval arithmetic in GP can greatly improve its generalisation performance. However, the training performance results suggest that adopting interval arithmetic invalidates more individuals during a run, which may reduce the effective size of the population and thus require larger populations and more generations to scale to harder problems. This remains an issue as long as interval arithmetic is solely used for static analysis within the the assignment of fitness to individuals. If the information provided by interval arithmetic could also be used within the search operators of GP, then this may result in fewer invalid individuals, and so we may recover some of the lost search performance.

Algorithm 1: `BuildTree`**: the safe initialisation tree generation algorithm.**

Input: *depth*: the current tree depth; *maxdepth*: the maximum required tree depth; F: a list of functions for inner nodes; T: a list of terminals (the input features of the problem); I: the known intervals of the terminal choices (input features, constants, …)

Output: A tree that produces output valid within the intervals defined by its subtrees

if `PickTerminal`($depth$, $maxdepth$, $|F|$, $|T|$) **then**
  $op \leftarrow$ random element from T
  return $\{root = op, left = \varnothing, right = \varnothing, interval = I[op]\}$
else
  $arg_0 \leftarrow$ `BuildTree`($depth + 1$, $maxdepth$, F, T, I)
  $ab \leftarrow arg_0[interval]$
  if `PickBinaryOperator`(F) **then**
    $arg_1 \leftarrow$ `BuildTree`($depth + 1$, $maxdepth$, F, T, I)
    $cd \leftarrow arg_1[interval]$
  else
    $arg_1 \leftarrow \varnothing$
    $cd \leftarrow$ [NaN, NaN]
  end if
  $op \leftarrow$ `SelectOperation`(F, ab, cd)
  $I_{op} \leftarrow$ `ComputeInterval`(op, ab, cd)
  return
  $\{root = op, left = arg_0, right = arg_1, interval = I_{op}\}$
end if

Algorithm 2: `CheckAndRepair`**: walks up the parent nodes of the tree to ensure that each point in the tree produces a valid interval.**

Input: *node* the parent node of the node that has just undergone a change (e.g., mutation or crossover); F: a list of functions for inner nodes;

Output: A tree that produces output valid within the intervals defined by its subtrees

$ab \leftarrow node[left][interval]$
if `BinaryOperator`($node[op]$) **then**
  $cd \leftarrow node[right][interval]$
else
  $cd \leftarrow$ [NaN, NaN]
end if
$I_{op} \leftarrow$ `ComputeInterval`($node[op]$, ab, cd)
if `InvalidInterval`(I_{op}) **then**
  shuffle F
  for $f \in F$ **do**
    $I_f \leftarrow$ `ComputeInterval`(f, ab, cd)
    if `ValidInterval`(I_f) **then**
      $node[op] \leftarrow f$
      $I_{op} \leftarrow I_f$
      break
    end if
  end for
end if
$node[interval] \leftarrow I_{op}$
`CheckAndRepair`($node[parent]$, F)

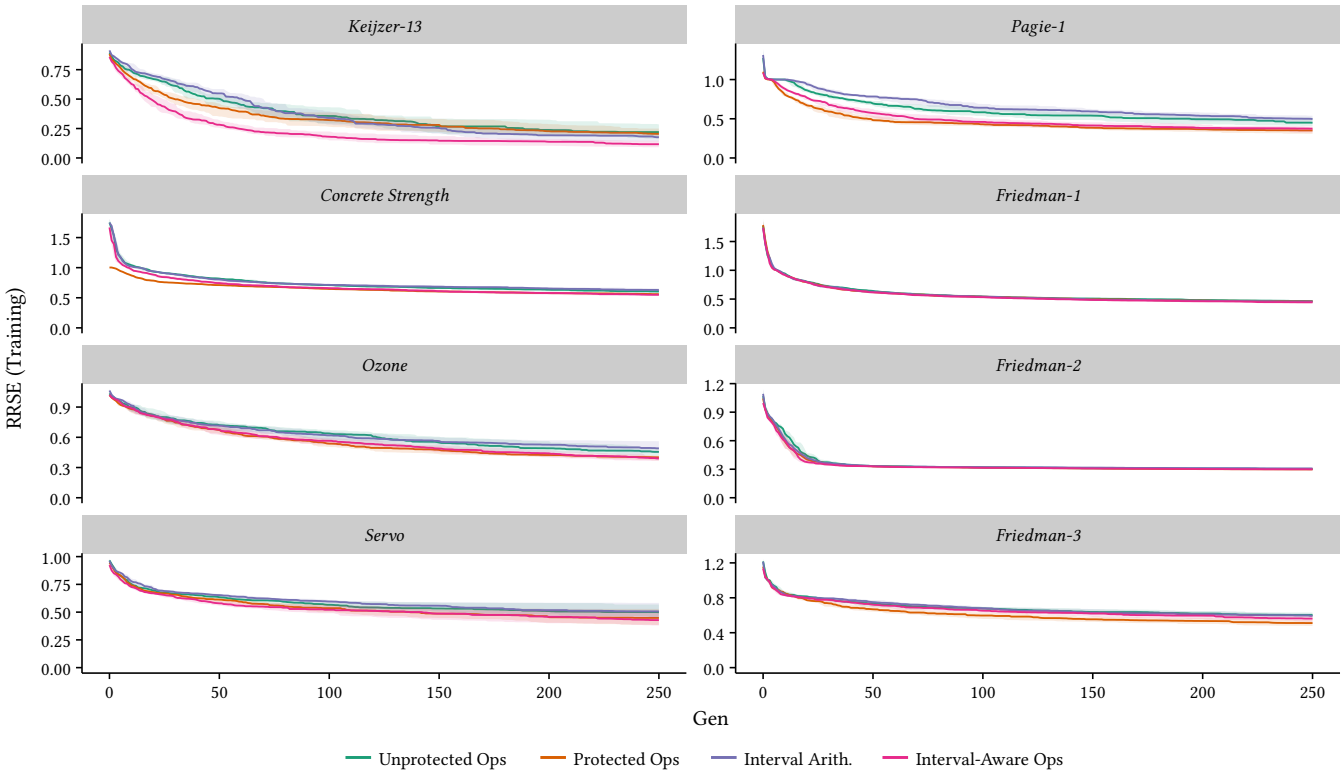

Figure 7: Training performance on selected benchmark problems.

Previous work identified that interval arithmetic could be used to preserve working bounds of solutions developed through geometric semantic genetic programming [5]. The so-called 'safe initialisation' method modified the tree initialisation process of GP such that the actual choice of operator at a given node was delayed until all the necessary child nodes had been created. Once subtrees were developed, their intervals could be computed and then safe initialisation could select an appropriate operator to maintain a valid interval for the entire tree. The end result is that invalid solutions are prevented from entering the population. When compared to an variant of GSGP that used logistic wrappers to manage intervals, safe-initialisation GSGP evolved at a much faster rate. Given that tree initialisation in GSGP is not substantially different from standard GP, it should be reasonably straightforward to integrate safe initialisation into standard GP. A definition of safe initialisation is given in Algorithm 1: this variant differs slightly from previous work by permitting both unary and binary operators (the previous definition allowed only binary operators).

Ensuring valid solutions during tree initialisation is only one aspect of operator protection that must be considered. In the previous section, it was shown that the proportion of invalid solutions within the population did not substantially change over time. This suggests that the search operators themselves may be contributing to the generation of invalid solutions. An example of how this may be happening is shown in Figure 6. Here mutation is applied to an individual with sane execution intervals. Likewise, the interval of the new mutant subtree is well defined. However, when this new tree is embedded into the solution, it interacts with its parent node to produce an undefined execution interval. A similar scenario can be imagined for crossover. What is needed, therefore, is a means by which the operation above the swapped subtree can be considered and, if necessary, repaired to produced valid intervals.

This paper extends the work done exploring safe initialisation — in addition to adopting safe initialisation in GP, we modify crossover and mutation operators so that they produce individuals that maintain useful intervals. Following crossover or mutation, the interval of the parent node of the modification site in the offspring is computed: if this interval is invalid, then a search is performed to find an operator that will take the child intervals as input and return a valid output interval. Once a valid operator is found, the check proceeds up the tree until the root node is encountered. If an operator cannot be found to produce a valid output interval, the offspring is flagged as producing an invalid output interval and assigned a low fitness. This repair mechanism is outlined in Algorithm 2.

5 EXPERIMENTAL COMPARISON

To test the performance of our interval preserving operators, we selected a range of problems from previous work. Two problems, *Keijzer-13* and *Pagie-1* are synthetic in nature — *Keijzer-13* has proven very difficult for standard GP in previous work, while the *Pagie-1* problem has been selected as its generating function:

$$f(x_1, x_2) = \frac{1}{1 + x_1^{-4}} + \frac{1}{1 + x_2^{-4}} \tag{4}$$

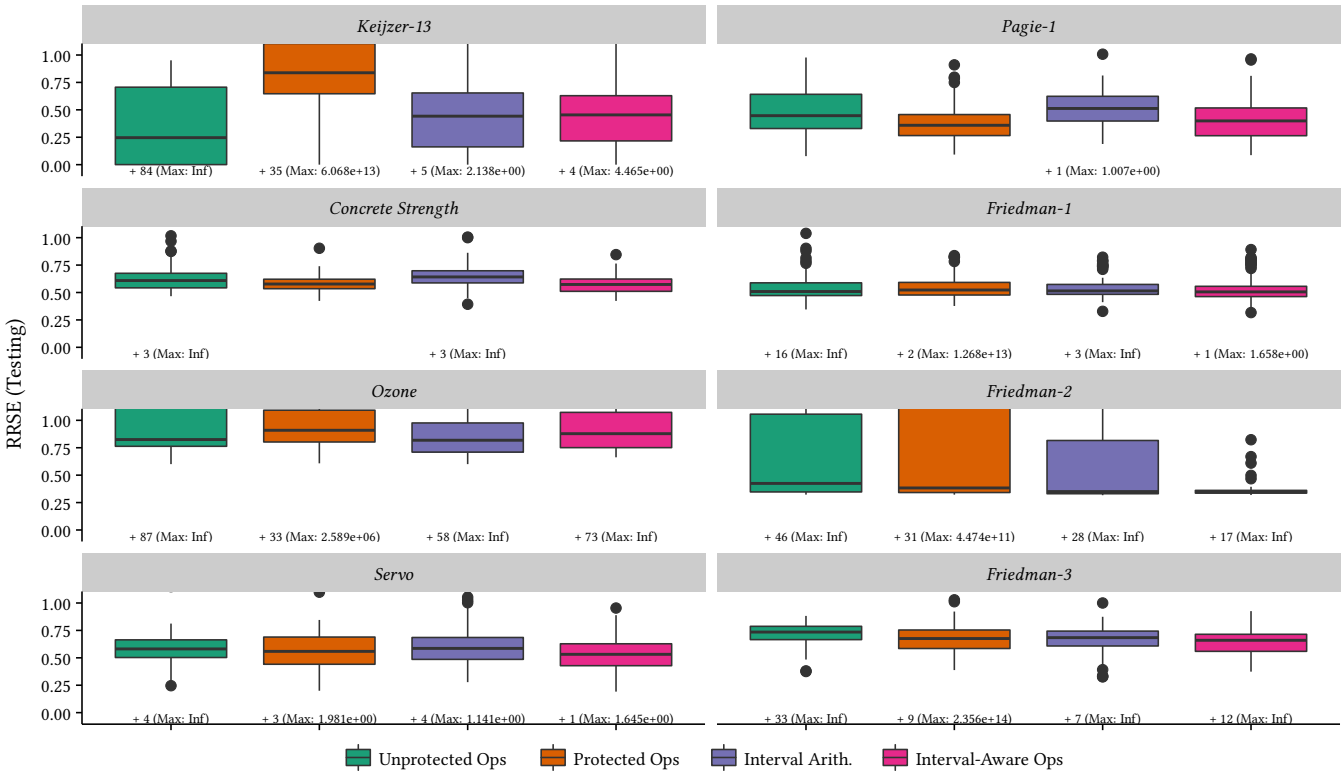

Figure 8: Test performance on selected benchmark problems.

has all its input features in the denominators, and their intervals include zero. Therefore, this should prove to be a difficult function for interval arithmetic. The other six functions have been used in previous work to test genetic programming and other machine learning methods [1, 2, 5, 7, 14]. All parameter settings used in these experiments are the same used earlier as outlined in Table 2. For the the *Keijzer-13* problem, we generated a mesh of points as outlined in previous work, and selected 20 points from this mesh for training for each run, with the remaining points used for testing. For the *Pagie-1* problem, we again used a mesh over the two variables to sample 676 points from the problem domain. We then sampled 68 points ($\tilde{1}0\%$) for training and used the rest for testing. For the remaining problems, we used 10 rounds of 10-fold cross validation to generate training and testing splits. Each combination of algorithm and problem was therefore run 100 times. For the *Keijzer-13* and *Pagie-1* the known intervals for the input variables was used: for the other problems, input intervals were estimated by examining the range of input variables in the training data. The impact of estimating input intervals will be examined later in this section.

The training performance of the various approaches on the eight problems is shown in Figure 7. As can be seen, for most of the problems, the revised search operators that consider the input intervals allow the GP system to evolve at a faster rate than when using interval arithmetic solely to evaluate individuals. With the exception of the *Concrete* and *Friedman-3* problems, this rate is comparable to the rate seen with protected operators. The performance of the interval-aware operators on *Pagie-1* is surprising and in contrast to simple interval arithmetic. As expected, using interval arithmetic solely for evaluation increases the difficulty, and search performance is restored once interval-aware operators are adopted.

The testing performance of the examined methods on the eight test problems is shown in Figure 8. The test performance is rather interesting. In particular, it is surprising that interval arithmetic, and interval-aware operators, despite having access to the input intervals, are not able to adequately control generalisation performance on three problems (*Ozone*, *Friedman-2* and *Friedman-3*). For the remaining problems, interval-aware operators generally present a small but useful improvement over straight interval arithmetic.

When looking globally, and considering only the median performance of the algorithms, there is some evidence that there is a difference between the methods, confirmed by a Friedman test where median algorithm performance is grouped by method and blocked by problem ($\chi^2 = 9.1519$, 3 degrees of freedom, p = 0.02734). However, looking at the boxplots suggests that, with the exception of the *Ozone* problem, methods including interval arithmetic are doing a better job of controlling destructive overfitting in most cases.

5.1 Sensitivity of Using Extimated Intervals

The performance of interval-based methods in the previous experiment was surprising. This is particularly so in the *Ozone* problem, where no method could adequately control the generalisation performance. Previous work suggests that interval arithmetic is

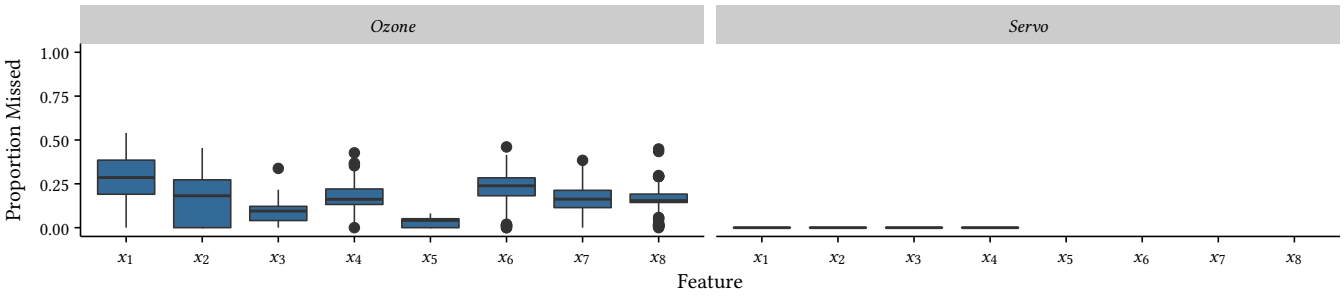

Figure 9: Misalignment between training and testing input intervals on the *Ozone* and *Servo* problems.

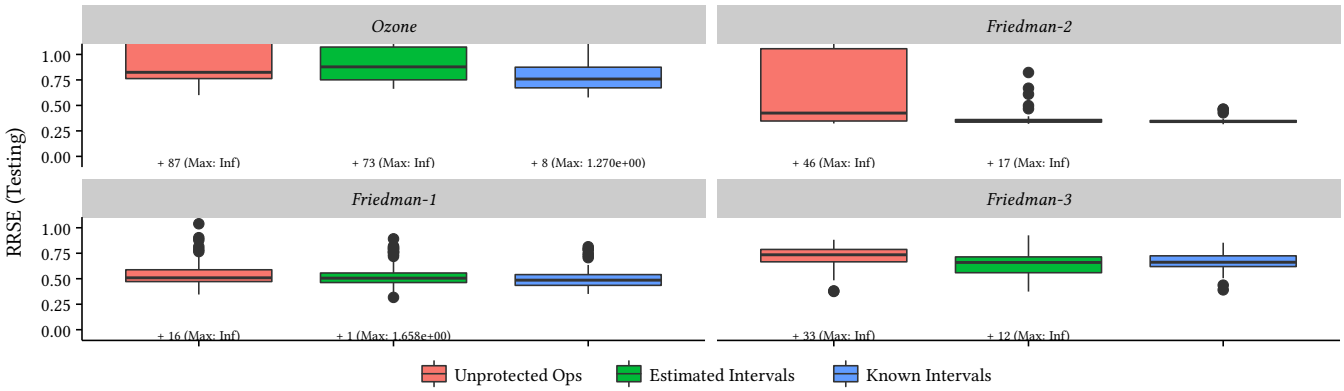

Figure 10: Training performance of interval arithmetic once measured (rather than estimated) intervals are provided.

guaranteed to provide correct interval information, but this naturally relies on the supplied intervals being an accurate reflection of what might be encountered outside of the training process. This suggests that the estimation process used to identify input intervals in the previous experiment is inadequate. Evidence towards supporting this is presented in Figure 9, which presents the proportion of the intervals of the input features that was not covered by the training data for the *Ozone* and *Servo* problems. These two problems present the two extremes of conditions encountered in the previous experiment. As can be seen, the intervals estimated from the training data in the *Ozone* problem are typically poor representations of the true intervals, whereas the intervals estimated within the *Servo* problem are well-aligned with the intervals that span testing.

To further examine the influence of the input intervals on algorithm performance, the four problems that were most problematic in the previous experiment were re-run with intervals that were either identified in previous work or measured from the entire data set rather than just the training data [7]. The results of this are shown in Figure 10. As can been seen, the performance of the interval-aware operators is greatly improved from that shown in Figure 8. In all cases, destructive overfitting has been eliminated — there are a handful of poor fitting models remaining for the *Ozone* problem, but these are due to poor search being performed in those runs, rather than through inadequate handling of intervals.

6 CONCLUSION AND FUTURE WORK

Symbolic regression has the potential to be a useful method for machine learing and data science. Traditionally, it has required the use of protected operators, and this has impacted on the quality of solutions evolved through GP. Previous work advocated the use of interval arithmetic to eliminate the need for protected operators. This paper extended previous work to uncover the dynamics of interval arithmetic during the evolutionary process. New operators were developed to greatly reduce the number of invalid solutions generated during a run, which allows evolution to proceed at a greater rate then by using interval arithmetic solely in the evaluation of individuals. Additionally, it was shown that operators using interval arithmetic are sensitive to the quality of the estimates of the input intervals — if intervals can be identified up front, then interval arithmetic is a reliable and safe way to conduct symbolic regression search. However, if there is poor alignment between the intervals used for training models and the intervals encountered subsequent to training, then interval arithmetic cannot greatly improve the generalisation performance of symbolic regression. Therefore, if interval arithmetic is to be used, then consideration must be made towards adequate identification of input intervals. In many cases, these intervals should not be difficult to establish — many measurements naturally lend themselves to reasonable identification of valid intervals (e.g., a person's age cannot be negative, and is not known to exceed 123 years). Therefore, the practical implications of this limitation of interval arithmetic are probably not as severe as demonstrated in this paper.

6.1 Future Work

The work presented in this paper opens several opportunities for future research. First, this paper only explored pure interval arithmetic for use in static analysis and operator execution. However, straight interval arithmetic is known to possess some issues in identifying true intervals. For example, if the interval of a variable x is [0, 1], then the interval of $x - x$ using straight interval arithmetic is [-1, 1], even though intuitively we know it is [0, 0]. Additionally, interval arithmetic is unable to identify correlations between variables, so the resulting intervals that it identifies are often much wider than would be encountered in practice. Affine arithmetic has been proposed as an alternative to interval arithmetic that considers variable correlations, and has been explored in the context of GP for static analysis [4, 13]. While we did not consider affine arithmetic in this work, it would be a useful exercise for future work to explore its use in the self-reparing operators we have developed.

Beyond a rudimentary analysis of the *Keijzer-10* problem, We have not performed a thorough analysis of the structure of the models that are produced with and without interval arithmetic. It would be interesting to compare the differences between these — it is likely that the search space is altered significantly through the integration of interval arithmetic, so it should make certain forms of operation (e.g., division) more difficult to identify. For example, the results on the *Pagie-1* problem using interval-aware operators are interesting and somewhat counter to what was expected. Interval arithmetic should find this a very difficult problem to search, as the input variables appear in the denominator of the problem, and the intervals associated with these inputs crosses zero (even if the input data never includes zero itself). However, despite this, the interval-preserving operators were able to perform at a level comparable to protected operators on this problem, suggesting that a suitable surrogate form was being evolved to replace the division.

Semantic operators are currently an active area of research in GP. There appears to be alignment between semantic methods and the interval-preserving operators in this work. Both attempt to idenfity properties of the subtrees in solutions and use these properties to guide the search process. More work exploring the alignment of these two concepts could potentially yield effective operators that allow GP to be applied to more difficult problems.

Finally, this paper used interval arithmetic for static analysis and repairing individuals corrupted by search operators. However, there are other ways that interval analysis could be used within GP that future work could explore. For example, previous work explored using the outputs of execution on data to identify areas in models that could be simplified into constant terms (numerical simplification) [10]. This allows aspects of trees to be simplified in a way that the alternative algebraic simplification cannot identify. Use of interval arithmetic (or, more likely affine arithmetic) within static analysis could potentially identify areas for simplification that would normally require the use of both algebraic and numerical simplication.

REFERENCES

[1] Alexandros Agapitos, James McDermott, Michael O'Neill, Ahmed Kattan, and Anthony Brabazon. 2014. *Higher Order Functions for Kernel Regression.* Springer Berlin Heidelberg, Berlin, Heidelberg, 1–12. DOI: http://dx.doi.org/10.1007/978-3-662-44303-3_1

[2] Leo Breiman. 2001. Random forests. *Machine learning* 45, 1 (2001), 5–32.

[3] Mauro Castelli, Sara Silva, and Leonardo Vanneschi. 2015. A C++ framework for geometric semantic genetic programming. *Genetic Programming and Evolvable Machines* 16, 1 (2015), 73–81.

[4] Luiz Henrique De Figueiredo and Jorge Stolfi. 2004. Affine arithmetic: concepts and applications. *Numerical Algorithms* 37, 1-4 (2004), 147–158.

[5] Grant Dick. 2015. Improving Geometric Semantic Genetic Programming with Safe Tree Initialisation. In *European Conference on Genetic Programming*. Springer International Publishing, 28–40.

[6] Grant Dick. 2017. Revisiting Interval Arithmetic for Regression Problems in Genetic Programming. In *Proceedings of the 2017 Annual Conference on Genetic and Evolutionary Computation*. ACM. To appear.

[7] Jerome H Friedman. 1991. Multivariate adaptive regression splines. *The annals of statistics* (1991), 1–67.

[8] Maarten Keijzer. 2003. Improving Symbolic Regression with Interval Arithmetic and Linear Scaling. In *Genetic Programming*, Conor Ryan, Terence Soule, Maarten Keijzer, Edward Tsang, Riccardo Poli, and Ernesto Costa (Eds.). Lecture Notes in Computer Science, Vol. 2610. Springer Berlin Heidelberg, 70–82.

[9] Maarten Keijzer, Martin Baptist, Vladan Babovic, and Javier Rodriguez Uthurburu. 2005. Determining Equations for Vegetation Induced Resistance Using Genetic Programming. In *Proceedings of the 7th Annual Conference on Genetic and Evolutionary Computation (GECCO '05)*. ACM, New York, NY, USA, 1999–2006. DOI:http://dx.doi.org/10.1145/1068009.1068343

[10] David Kinzett, Mark Johnston, and Mengjie Zhang. 2009. Numerical simplification for bloat control and analysis of building blocks in genetic programming. *Evolutionary Intelligence* 2, 4 (2009), 151–168.

[11] Mark Kotanchek, Guido Smits, and Ekaterina Vladislavleva. 2008. *Trustable symbolic regression models: using ensembles, interval arithmetic and pareto fronts to develop robust and trust-aware models.* Springer US, Boston, MA, 201–220. DOI: http://dx.doi.org/10.1007/978-0-387-76308-8_12

[12] John R. Koza. 1992. *Genetic Programming: On the Programming of Computers by Means of Natural Selection.* MIT Press, Cambridge, MA, USA.

[13] Cassio L Pennachin, Moshe Looks, and João A de Vasconcelos. 2010. Robust symbolic regression with affine arithmetic. In *Proceedings of the 12th annual conference on Genetic and evolutionary computation*. ACM, 917–924.

[14] J Ross Quinlan. 1993. Combining instance-based and model-based learning. In *Proceedings of the Tenth International Conference on Machine Learning*. 236–243.

[15] Cao Truong Tran, Mengjie Zhang, Peter Andreae, and Bing Xue. 2016. Directly Constructing Multiple Features for Classification with Missing Data Using Genetic Programming with Interval Functions. In *Proceedings of the 2016 on Genetic and Evolutionary Computation Conference Companion (GECCO '16 Companion)*. ACM, New York, NY, USA, 69–70. DOI: http://dx.doi.org/10.1145/2908961.2909002

[16] Grégory Valigiani, Cyril Fonlupt, and Pierre Collet. 2004. *Analysis of GP Improvement Techniques over the Real-World Inverse Problem of Ocean Color.* Springer Berlin Heidelberg, Berlin, Heidelberg, 174–186. DOI: http://dx.doi.org/10.1007/978-3-540-24650-3_16

[17] David R. White, James McDermott, Mauro Castelli, Luca Manzoni, Brian W. Goldman, Gabriel Kronberger, Wojciech Jaśkowski, Una-May O'Reilly, and Sean Luke. 2013. Better GP benchmarks: community survey results and proposals. *Genetic Programming and Evolvable Machines* 14, 1 (2013), 3–29. DOI: http://dx.doi.org/10.1007/s10710-012-9177-2